

Detection and Modelling of non-Gaussianity in MR Diffusion Imaging

D. C. Alexander¹, G. J. Barker², S. R. Arridge¹.

¹ *Dept. Computer Science, UCL (University College London), Gower St., London, WC1E 6BT,*

UK.

² *Institute of Neurology, UCL, Queen Sq., London, WC1N 3BG, UK.*

Running head: **non-Gaussian ADC Profiles.**

Contact author:

Daniel Alexander

Dept. Computer Science, UCL (University College London), Gower Street, London, WC1E
6BT. UK

Tel.: +44 20 679 2419

Fax: +44 20 7387 1397

Email: Daniel.Alexander@cs.ucl.ac.uk

Abstract

This paper details the observation of non-Gaussian apparent diffusion coefficient (ADC) profiles in multi-direction, diffusion-weighted MR data acquired with standard clinical imaging parameters ($b \approx 1000 \text{ s/mm}^2$). A technique is described for modelling the profile of the ADC over the sphere, which handles the higher order, non-Gaussian effects that can occur at, for example, intersections of different tissue types or white matter fibre tracts. When these higher order effects are significant, the common diffusion tensor model for the ADC profile is inappropriate, since it is based on the assumption of an underlying Gaussian diffusion process. A sequence of models of increasing complexity is obtained by truncating the spherical harmonic expansion of the ADC measurements at several orders. Further, a method is described for selection of the most appropriate of these models, in order to describe the data adequately and without overfitting. The combined procedure can be used to classify the profile at each voxel as isotropic, anisotropic Gaussian or non-Gaussian. We use it to show that higher order profiles arise consistently, in various regions of the human brain where complex tissue structure is known to exist, and can be observed in data typical of clinical scanners. The performance of the procedure developed is characterised carefully using synthetic data in order to demonstrate that the observed effects are genuine, which validates use of the method as an indicator of pathology that affects tissue structure.

Keywords: Diffusion tensor magnetic resonance imaging, human brain, non-Gaussian diffusion, spherical harmonic.

1 Introduction

Diffusion tensor magnetic resonance imaging (DT-MRI) [1] has become popular, because of the insight it provides into the structural connectivity of tissue [2, 3]. DT-MRI provides an image of the *apparent diffusion tensor* (ADT) within each voxel. The ADT is a model of the *apparent self-diffusion coefficient* (ADC) profile and provides an estimate of the ADC in every direction. The *self-diffusion coefficient* (SDC) [4] is defined at a point in space and is a function of direction. In a particular direction, the SDC is proportional to the mean square displacement of molecules over unit time. The ADC is a spatial average of the SDC over some volume, e.g. an image voxel, which can be measured by *diffusion weighted magnetic resonance imaging* (DW-MRI). In the simplest form of quantitative DW-MRI, two measurements are acquired at each voxel: a DW measurement, $S(\hat{\underline{x}})$, acquired in the presence of a DW gradient in direction, $\hat{\underline{x}}$, together with an unweighted measurement, S_0 . These two measurements are related by the ADC in direction $\hat{\underline{x}}$, $d(\hat{\underline{x}})$, through the following equation [5, 6]:

$$S(\hat{\underline{x}}) = S_0 e^{-bd(\hat{\underline{x}})}, \quad (1)$$

where b is the diffusion weighting factor. (1) can be rearranged to give an expression for $d(\hat{\underline{x}})$:

$$d(\hat{\underline{x}}) = \frac{1}{b} \log \left(\frac{S_0}{S(\hat{\underline{x}})} \right). \quad (2)$$

If molecular diffusion is governed by the same Gaussian process everywhere within an image voxel, then the 3D displacement, \underline{x} , of a molecule after fixed time t is drawn from a multivariate Gaussian distribution [4] with density function

$$p_t(\underline{x}) = \frac{1}{\sqrt{(4\pi t)^3 |D|}} \exp\left[\frac{-(\underline{x} - \underline{x}_0)^T D^{-1} (\underline{x} - \underline{x}_0)}{4t}\right]. \quad (3)$$

It is clear from (3) that in this Gaussian case, $p_t(\underline{x})$ is completely defined by the covariance matrix, D , which is a 3x3 symmetric tensor called the *diffusion tensor*. In this case, D also relates directly to the ADC profile [1, 4]:

$$d(\hat{\underline{x}}) = \hat{\underline{x}}^T D \hat{\underline{x}}. \quad (4)$$

Thus,

$$S(\hat{\underline{x}}) = S_0 e^{-b \hat{\underline{x}}^T D \hat{\underline{x}}}. \quad (5)$$

In general [10, 11], $p_t(\underline{x})$ has a Fourier relationship with the spin echo attenuation, $S(\hat{\underline{x}})/S_0$, measured with varying DW gradient strength, the square of which is proportional to the b-factor. It is straightforward to show that this relationship holds between equations (3) and (5) in the Gaussian case [11].

It has long been recognised [1, 11-16], however, that the Gaussian, diffusion tensor model is inappropriate when complex tissue structure is found within a single image voxel. There are alternative models that can capture certain non-Gaussian effects that occur in these circumstances. The simplest alternative is the multi-tensor model [12, 13]. This model is based on the assumption that a voxel contains n different tissue types each in proportion p_i ($\sum_i p_i = 1$, $i = 1, \dots, n$) and with Gaussian diffusion characteristics described by DT , D_i , and further that there is no exchange of molecules between these separate compartments. $p_t(\underline{x})$ then becomes a weighted sum of Gaussians and (5) becomes:

$$S(\hat{\underline{x}}) = S_0 \sum_{i=1}^n p_i e^{-b \hat{\underline{x}}^T D_i \hat{\underline{x}}}. \quad (6)$$

With this model, the ADC profile can have shape very different to that described by (4), which is often modelled poorly by a single DT [12]. Figure 1 shows ADC profiles simulated from prolate, oblate and isotropic Gaussian diffusion processes, together with profiles obtained from bi-tensor models of mixtures of these processes. Note the characteristic peanut shape of the prolate ADC profiles and the filled bagel or red blood cell shape of the oblate profile. The corresponding $p_i(\underline{x})$'s, have the more familiar cigar shaped, in the prolate case, and Frisbee shaped, in the oblate case, ellipsoidal contours.

In DT-MRI, the diffusion is assumed to be Gaussian and the DT, D , is estimated from samples of the ADC. A minimum of six measurements of the ADC are required to fit the six parameters of the DT. However, a more robust approach [7] is to oversample the profile, by acquiring many more DW images in a large number of directions spread evenly over the unit hemisphere (antipodal symmetry of $p_i(\underline{x})$ and hence the ADC profile is assumed). Such a scheme affords the possibility of observing more complex, non-Gaussian ADC profiles should they arise.

Alexander, et al [12] analysed the behaviour of the ADC profile within voxels containing multiple tissue compartments. They showed how the observability of higher order profiles increases (their non-Gaussian shape becomes more pronounced) with the size of the diffusion weighting factor, b in Eqs (1) and (2). They used the instability of the DT and its derived scalar measures such as the mean diffusivity and fractional anisotropy to locate regions in the brain where the DT model of the ADC profile is poor. Several regions in data acquired with $b = 3000$ s/mm² were highlighted through this approach, including the pons, the corpus callosum, the cingulum, the internal capsule and the arcuate fasciculus.

Frank [13] showed that a 4th order SH series provides a first approximation to the ADC profile obtained from a multi-tensor diffusion process. Frank fitted a 4th order SH series to ADC data

acquired with $b = 3000 \text{ s/mm}^2$ and showed that significant 4th order (i.e. non-Gaussian) components arise in similar locations of the human brain to those highlighted in [12] (see above).

In other related work, Wedeen and Tuch, et al [14, 15] used q-space techniques to measure $p_t(\underline{x})$. This technique exploits the Fourier relationship between $p_t(\underline{x})$ and the spin echo attenuation by acquiring a large number of measurements of the latter over a wide range of b-values in order to obtain enough samples of the FT of $p_t(\underline{x})$ to perform a stable inversion. Distinctly non-Gaussian $p_t(\underline{x})$'s were observed in both the human brain and heart particularly at locations where anisotropic fibres cross within a single voxel.

In this work, we investigate the observability of higher order ADC profiles in DW data acquired using acquisition parameters more typical of those used clinically. Our sequence consists of a multiple gradient direction scheme based on the work of Jones, et al [7] using a 1.5T scanner, with a maximum b-value of approximately 1000 s/mm^2 . We use the spherical harmonic (SH) series to provide a hierarchy of models for the ADC profile. In section **Error! Reference source not found.**, an efficient and robust method for fitting the SH series to DW-MR data is described together with a method, based on the analysis of variance (ANOVA) test for addition/deletion of variables, for selecting the most appropriate level of truncation of the series. The combined fitting and model selection procedures developed can be used to classify the profile in each voxel as isotropic, anisotropic Gaussian or non-Gaussian and thus to produce maps of where these different types of behaviour occur. These maps should provide extra diagnostic information in pathologies involving neuronal loss, degeneration or demyelination, since higher order behaviour will tend to disappear in the affected areas. We apply the method to both in vivo human brain data and to synthetic data for performance evaluation and validation. These experiments and their results are presented in section 3 and we conclude in section 4.

2 Methods

2.1 Models for the ADC Profile

In this section, we describe how the SH series can be used to model the ADC profile. We define the SH series, describe how its coefficients are computed for a given function or set of sampled data and discuss the relationship to more familiar models of the ADC.

2.1.1 The Spherical Harmonic Series

The SH series [18] is analogous to the rectilinear Fourier series and provides an orthonormal basis of functions on the sphere that can be combined linearly to represent any complex valued spherical function. It consists of a set of functions $Y_{l, m}: S_2 \rightarrow \mathbf{C}$, where S_2 is the unit sphere in 3D, which we parametrise by $\theta \in [0, \pi)$ and $\varphi \in [0, 2\pi)$, the angles of colatitude and longitude, respectively; \mathbf{C} is the set of complex numbers. $l = 0, 1, 2, \dots$ defines the *order* of the SH and $m \in \{-l, \dots, 0, \dots, l\}$ indexes the $2l+1$ SH functions of order l . The individual SH functions are defined in terms of the *associated Legendre polynomials*, see [18, 19].

Any spherical function, $f: S_2 \rightarrow \mathbf{C}$, can be written as a linear combination of the SHs:

$$f(\theta, \varphi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l c_{l,m} Y_{l,m}(\theta, \varphi). \quad (7)$$

The complex coefficients, $c_{l,m}$, of the spherical harmonics in (7) are given by:

$$c_{l,m} = \int_0^{2\pi} \int_0^{\pi} f(\theta, \varphi) Y_{l,m}^*(\theta, \varphi) \sin \theta d\theta d\varphi. \quad (8)$$

In (8), and henceforward, $*$ denotes the complex conjugate.

2.1.2 Fitting the Series to Sampled Data

In the discrete case, where we have sampled data, $\underline{\mathbf{F}} = \{f(\theta_i, \varphi_i), i = 1, \dots, n\}$, one way to proceed is to replace the integral in (8) by a summation and approximate the $c_{l, m}$ by

$$c_{l, m} \approx \frac{4\pi}{n} \sum_{i=1}^n f(\theta_i, \varphi_i) Y_{l, m}^*(\theta_i, \varphi_i). \quad (9)$$

As noted by Brechbuhler, et al [20] however, the use of (9) directly to compute the $c_{l, m}$ often provides poor estimates for their values and a more robust approach is to compute the least squares solution. We adopt a similar approach here. First, we define an enumeration of the SH series, so that each SH is indexed by a single, unique integer, $j(l, m) = l^2 + l + m$. We select a maximum order for the series, l_{\max} , and then define a $n \times j(l_{\max}, l_{\max})$ complex matrix, \mathbf{X} , to be the matrix with elements $X_{i, j(l, m)} = Y_{l, m}(\theta_i, \varphi_i)$. If we define $\underline{\mathbf{C}}$ to be the $j(l_{\max}, l_{\max})$ component vector of SH coefficients and $\underline{\mathbf{F}}$ is the n -component vector of sample points, defined above, then [20]

$$\underline{\mathbf{C}} = \mathbf{M} \underline{\mathbf{F}}, \text{ where } \mathbf{M} = (\mathbf{X}^* \mathbf{T} \mathbf{X})^{-1} \mathbf{X}^* \mathbf{T}. \quad (10)$$

l_{\max} can be chosen by consideration of the number of free parameters defining the series up to each order, which should be less than or equal to the number of sampled points.

2.1.3 Modelling DW-MR data

In its most general form, the SH series can represent any complex-valued function of the sphere in 3D. The set of functions required to model ADC profiles, however, is more constrained. In particular, the ADC is real-valued and exhibits antipodal symmetry, so that $D(\theta, \varphi) = D(-\theta, -\varphi)$ for all θ, φ .

The real valued constraint ensures both that the imaginary part of $c_{l, 0}$ is zero for all l , and that the following relationships exist between the coefficients of $Y_{l, m}$ and $Y_{l, -m}$:

$$c_{l,m} = (-1)^m c_{l,-m}^* \quad (11)$$

The SHs of odd order all define asymmetric functions on the sphere, whereas the even orders are all symmetric. The antipodal symmetry of the ADC ensures that it can be represented by a series consisting only of even order SHs.

These constraints dramatically reduce the number of parameters required to describe SH models truncated at each order. The numbers of free parameters required for SH series of each order, for both general spherical functions and real, symmetric functions, are listed in Table 1.

Order of SH model	No. of SH functions at this order.	Total no. of SH functions up to this order.	No. parameters for general functions.	No. parameters for real, symmetric functions.
0	1	1	2	1
1	3	4	8	1
2	5	9	18	6
3	7	16	32	6
4	9	25	50	15
6	13	49	98	28
8	17	81	162	45
n	2n + 1	(n+1) ²	2(n+1) ²	(n+1)(n/2+1)

Table 1 Number of parameters defining SH models up to different orders.

Eq. (10) tells us that the calculation of the SH coefficients from the set of sampled points on the ADC profile is a linear transformation and so can be computed very efficiently. Moreover, the set of sampled points at each voxel in the image correspond to the same set of directions, so that $\{(\theta_i, \phi_i), i = 1, \dots, n\}$ is fixed over the image. The matrix M , of Eq. (10), therefore need only be calculated once in order to compute \underline{C} at each voxel via the linear transformation described in (10).

2.1.4 Models at Different Orders

If the SH series is truncated at order 0, the series provides an isotropic model, since the single 0th order SH is constant over the sphere. If the series is truncated at order 2, it provides a model that is completely equivalent to the familiar DT model. In fact, it is straightforward to obtain an expression for the DT in terms of the 0th and 2nd order SH coefficients, by equating the right hand sides of Eqs. (4) and (7). The resulting expression (omitted for brevity) demonstrates that the DT is also a linear transformation of the ADC samples and fitting the SH series up to order 2 provides an alternative way to fit the DT. When we include higher order members of the series, a range of more complex shapes becomes available. In particular, at order 4 we can obtain models with two pairs of peaks that approximate the shapes obtained from the bi-tensor model, such as those shown in Figure 1.

2.2 Model Selection

Once we have computed the coefficients of the SH series for a given set of sample points, a hierarchy of increasingly complex models is obtained by truncating the series at each order, $l = 0, 2, 4, \dots, l_{\max}$ (the coefficients of the SHs above order l are set to zero). Although higher order models are more descriptive, in many cases they may not be necessary to describe the underlying function from which the measurements were taken. For example, if the underlying function is isotropic, we only need a series up to order 0 to describe it and higher order terms will only represent noise added to the data during the imaging process. If the underlying function is described by a DT, we only need a series up to order 2. The methodology for selecting the appropriate level of truncation of the series is described in this section.

We use ANOVA to determine whether the addition of more parameters to the model, i.e., truncating the series at a higher order as opposed to a lower order, significantly changes the fit of

the model to the data [21]. Given a set of N sampled points, together with a lower order model M_1 with p_1 free parameters and a higher order model M_2 with p_2 free parameters, the appropriate statistic for the F-test of the hypothesis that the two models are equivalent, i.e., the lower order model is sufficient, is

$$F(M_1, M_2) = \frac{(N - p_2 - 1)(\text{Var}(M_2) - \text{Var}(M_1))}{(p_2 - p_1)\text{MSE}(M_2)}. \quad (12)$$

The degrees of freedom are $N - p_2 - 1$ and $p_2 - p_1$ [21]. In (12), $\text{Var}(M)$ denotes the variance of model M about its mean value and $\text{MSE}(M)$ denotes the mean squared error between model M and the N sampled points.

The full algorithm for modelling the ADC profile in one voxel is outlined below:

- Compute the coefficients of the even SH series up to order l_{\max} using (10).
- Truncate the series at order 0 to obtain model M_0 .
- Set $i = 2$ and $a = 0$.
- While ($i \leq l_{\max}$),
 - Truncate the series at order i to obtain model M_i .
 - Compute $F(M_a, M_i)$ using equation (12).
 - Define the null hypothesis to be that models M_i and M_a are equivalent.
 - Compute the critical value, T , for F such that if $F < T$ then the probability of the null hypothesis is less than a decision threshold, α .
 - If $F < T$ (null hypothesis is rejected)
 - Set $a = i$.
 - Set $i = i + 2$.
- Return model M_a .

3 Experiments

The central hypothesis to be tested here is that ADC profiles corresponding to non-Gaussian diffusion processes exist in the human brain and can be observed in data collected on a standard

clinical MR scanner using the methods described in the previous section. In order to verify this hypothesis we apply our methods to a variety of synthetic data and to human brain DW-MR data acquired from a such a scanner.

The synthetic data is created by a Monte-Carlo simulation of the imaging process. This data is first used to set the parameters of the model selection procedure and subsequently to characterise its performance in terms of the number and type of misclassifications that occur for noisy profiles of known order. We go on to show the results of applying our procedure to human brain data, which exhibit clusters of higher order models in various well-defined regions. In order to confirm that the higher order regions observed in the human brain data are genuine, we perform a final simulation in which the Gaussian (order 2 SH) models found in these regions are emulated and run through our procedure.

3.1 *Simulation Experiments*

In this section, we test our procedure on various synthetic ADC profiles, corresponding to different types of tissue found in the brain. We simulate the imaging process used to acquire the human brain data described in the next section in order to generate noisy measurements of particular ADC profiles. A range of DTs is used to provide models for profiles corresponding to isotropic and anisotropic Gaussian diffusion and we use a bi-tensor model, as described in the introduction (see Eq. (6)), to obtain ADC profiles corresponding to mixed tissue and fibre crossings.

For a given profile, we simulate a 128x128 2D array of noisy measurements of the profile. In order to ensure that the noise is realistic [22], Gaussian noise is added in the time/frequency domain at a level corresponding to that observed in the scanner in the absence of any signal. Noisy images are obtained in this way, which correspond to each of the 60 DW images acquired

in our sequence (see next section) together with 3 unweighted images. ADC measurements in each of the 60 sampled directions are then reconstructed and our procedure is applied.

The first set of experiments performed are designed to provide appropriate settings for the decision threshold α , used in the F-test in the algorithm described in section 2.2. Lower settings of α cause the null hypotheses to be more difficult to reject, which will generally result in an increase in the proportions of lower order models. Higher values of α result in greater numbers of higher order models and in general there is a trade off between the number of overfitted and underfitted voxels, the occurrence of both of which we would like to minimise. We simulate the following Gaussian ADC profiles for testing [3]:

- Grey matter (GM), where the diffusion is approximately isotropic with eigenvalues [700, 700, 700] ($\times 10^{-6}$ mm²/s) and the unweighted signal is chosen so that its signal to noise ratio, $\sigma_0 = 35$, which is typical in our scanner data.
- Cerebro-spinal fluid (CSF), eigenvalues [3000, 3000, 3000] ($\times 10^{-6}$ mm²/s) and $\sigma_0 = 115$.
- Prolate white matter (WM), eigenvalues [1700, 200, 200] ($\times 10^{-6}$ mm²/s), $\sigma_0 = 35$.
- Oblate WM, eigenvalues [950, 950, 200] ($\times 10^{-6}$ mm²/s), $\sigma_0 = 35$.

We also simulate profiles from orthogonal crossing fibres in equal proportions, where each fibre is represented by a prolate DT with eigenvalues [1700, 200, 200] ($\times 10^{-6}$ mm²/s) and $\sigma_0 = 35$.

For each profile, α is varied and the number of voxels at which each model order is selected is recorded. The proportions of each model type selected for each profile are plotted in Figure 2. From Figure 2 we see that the most appropriate setting for the decision threshold for discriminating effectively between order 0 and order 2 profiles is very low – around 10^{-20} , however, for discriminating effectively between order 2 and order 4 models, the decision threshold should be set higher – in the region of 10^{-7} . Thus, in the algorithm outlined in section

2.2, we allow the decision threshold to vary depending on the order of the current best model, M_a , and choose $\alpha = 10^{-20}$ if the current best model is order 0 and $\alpha = 10^{-7}$ otherwise.

Once appropriate decision thresholds for our F-tests have been chosen, we can characterise the performance of our procedure on a wider range of input data, which enables us to interpret results obtained from real data reliably. A series of oblate and prolate Gaussian ADC profiles are tested in which the anisotropy is increased from zero to approximately the highest levels observed in the brain. The classification rates of our procedure applied to these data are plotted in Figure 3 (prolate profiles) and Figure 4 (oblate profiles).

We also simulate profiles from crossing fibres intersecting at angles of 90, 67.5, 45, and 22.5 degrees, and in proportions 1:1 and 3:1. Each fibre is represented by a prolate DT with eigenvalues [1700, 200, 200] ($\times 10^{-6}$ mm²/s) and $\sigma_0 = 35$. Classification rates from these profiles are plotted in Figure 5. Profiles obtained from partial volumes of CSF and prolate WM together with those from GM and prolate WM are also simulated and their classification rates are plotted in Figure 6. Values of σ_0 for mixed CSF and WM are obtained by combining the pure values of 115 and 35 in corresponding linear proportions.

Figure 3 and Figure 4 demonstrate that, with the selected thresholds, our procedure will correctly identify anisotropic, Gaussian profiles with good consistency, when the difference between the largest and smallest eigenvalues of the DT is approximately 300 ($\times 10^{-6}$ mm²/s) or greater. At this lowest level of detected anisotropy, the rate of misclassification at order 4 is around 2%, but the rate increases steadily with increasing anisotropy to about 8% for the most anisotropic DT. Misclassification at orders above 4 is negligible. These figures also demonstrate the classification of isotropic GM profiles, which is reliably order 0 – the rate of misclassification at order 2 is less than 0.1%. These rates are consistent for other isotropic DTs

(with larger eigenvalues), in particular, the classification rates for simulated isotropic CSF (shown in Figure 6) are very similar. Figure 5 shows that our simulated fibre intersection profiles are classified as order 4 with good reliability when the fibres are in equal proportion and are orthogonal, when only 3% of the profiles are underfitted as order 2 profiles. However, the rate of misclassification at order 2 increases significantly as the proportion of the two fibres becomes less balanced and as the two fibres become more parallel. This is reasonable to expect, because both these effects cause the deviation of the profile from Gaussian to decrease. Misclassification above order 4 for all these bi-tensor profiles varies between 0.5% and 1%. The results in Figure 6 demonstrate that the deviation from Gaussian behaviour that is obtained by mixing isotropic and anisotropic Gaussian diffusion with the multi-tensor model is not reliably detected by our procedure with these imaging parameters. Only a minor increase in the order 4 classification rate is observed for these mixed profiles.

3.2 Human Data Experiments

DW-MR data was acquired from four healthy volunteers using a protocol similar to that outlined by Jones [7]. All subjects were scanned with the approval of the joint National Hospital and Institute of Neurology ethics committee and gave informed, written consent. Three unweighted ($b = 0 \text{ s/mm}^2$) images were acquired together with 60 DW images with different gradient directions spread evenly over the sphere, with $b \approx 1000 \text{ s/mm}^2$ in each case. The image array is 128×128 in plane with a field of view of 220mm and a total of 42 slices evenly spaced at 2.5mm were acquired.

The procedures described in section 2 were applied to this clinical data. Since 60 samples are available for each ADC profile, the maximum series order that we can fit is 8. We thus compute the coefficients of all even order SHs up to and including order 8 and then apply the model

selection procedure described in section 2.2, with the decision thresholds selected in the previous section, to choose the appropriate level of truncation of the series. Prior to model selection, two thresholds on σ_0 are applied: if $\sigma_0 < 8.5$, the voxel is classified as background and no model is assigned; if $\sigma_0 > 85$, we assume the voxel corresponds to CSF and assign an order 0 model, since flow artefacts in these regions make the measurements unreliable.

Figure 7 shows maps of the order chosen by our model selection procedure (left) for three axial slices of one of our data sets, together with mean diffusivity (right) and anisotropy weighted colour maps indicating the principal direction of the DT at each point (middle). In the colour images, red, green and blue intensities correspond to the size of the x, y and z components of the DT principal direction vector, weighted by the fractional anisotropy, [2], as proposed by Pajevic and Pierpaoli, [23].

A detailed anatomical analysis of the results is beyond the scope of this paper, but we will highlight some interesting features. Firstly, our procedure appears to separate regions of isotropic and anisotropic diffusion successfully. Regions of grey matter are mostly assigned order zero models, while regions of white matter, both dense and peripheral, are mostly assigned order 2. A significant percentage of voxels are assigned higher order models. The spatial distribution of these voxels exhibits distinct clusters in the image, which have clear symmetry about the midline of the brain. This clustering and symmetry strongly suggests that the higher order models correspond to genuine anatomical effects.

Each slice shown in Figure 7 contains an anatomic region that contains clusters of higher order voxels consistently in each of our four data sets. In the top slice, the region corresponding to the pons contains a dense cluster of order 4 voxels (label 1). In this region, the right-left trans-ponsine tracts cross the inferior-superior pyramidal tracts, causing partial volume effects between

these two orthogonal fibres. In the middle slice, dense clusters of order 4 voxels are found in the optic radiation on both sides of the brain (label 2). These occur precisely at the point where the anterior-posterior tracts of the optic radiation cross the predominantly right-left fibres of the corpus callosum. In the bottom slice, large clusters of order 4 voxels can be observed within the corona radiata (label 3). Again this is reasonable to expect, since complex interactions between the diverging fibres of the corona radiation and the U-fibres occur in these areas.

Figure 9 shows typical ADC profiles from each of these three regions, together with the order 0, 2, 4, 6 and 8 spherical harmonic models. In each case, it is clear that there is significant difference between the order 4 and order 2 models, which indicates significant non-Gaussian behaviour. The models with order greater than 4 do not appear to change the overall profile shape significantly further and serve only to incorporate noise effects. The measurement from the pons is particularly striking and is typical of that predicted by the bi-tensor model for two orthogonally intersecting WM fibres, see Figure 1. The measurement from the optic radiation is similar but oriented within the axial plane and somewhat less pronounced, possibly due to a less balanced mix between the two fibres or a smaller angle of intersection. The profile from the corona radiata is more difficult to interpret and may arise from more complex interactions that are not modelled naturally by a pair of DTs.

We note that the proportions (excluding background voxels) of each model order that were assigned over the whole of the data set shown in Figure 7 were: order 0 – 53%, order 2 – 43%, order 4 – 4%, order 6 – 0.05% and order 8 – 0%. These proportions are similar in the other three data sets tested. Orders maps for similar slices to those shown in Figure 7 for the other three data sets are shown in Figure 8.

3.3 Further Synthetic Experiments

In this section we perform some additional synthetic experiments to verify that the observations made on the human brain data are due to genuine higher order effects rather than artefacts of our fitting and model selection procedures. Since the set of DTs tested in the synthetic experiments described in section 3.1 is not exhaustive, it is possible that the regions we observe in the human brain data simply exhibit Gaussian diffusion processes that are particularly prone to overfitting. Thus, for completeness, we include this additional set of experiments to test this assertion. In each region, we fit the DT at each voxel and model the distribution of DT eigenvalues, $\lambda_1 > \lambda_2 > \lambda_3$, together with the value of σ_0 . We then use a Monte-Carlo method to draw samples from this distribution and test the likelihood that the corresponding Gaussian ADC profiles are overfitted with higher order ADC profile models. σ_0 is included in this model, because the noise levels in ADC measurements depend on σ_0 as well as the magnitude of the ADC itself, [24].

We concentrate on three of the regions highlighted in the previous section: the pons, the optic radiation and the corona radiata. For each of these anatomic regions, a sub-region of the order 4 cluster was selected and a model of the distribution of DT shapes was obtained by fitting a 4D multivariate Gaussian model to the distribution of vectors $(\lambda_1, \lambda_2, \lambda_3, \sigma_0)$ within that sub-region. A similar procedure to that outlined at the start of section 3.1 was used to generate arrays of noisy ADC measurements corresponding to DTs and σ_0 values drawn from these distributions. The fitting and model selection procedure was then applied to these profiles and the classification rates are shown in Figure 10, in which we see that the rate of classification above order 2 is less than 5% in each case. These results verify that if the underlying diffusion process had been Gaussian, our procedure would have been effective and assigned most profiles in these

regions order 2 models. Accordingly, we conclude that we are observing genuine higher order effects in these regions.

4 Discussion

We have described methods for modelling and detection of higher order ADC (apparent diffusion coefficient) profiles and shown that such profiles can be observed in standard clinical DW-MR data. The SH (spherical harmonic) series up to order 8 was fit to samples of the ADC profile in each voxel, which provides a sequence of models of increasing complexity. A series of ANOVA (analysis of variance) tests was used to find the simplest of these models that adequately describes the data. The performance of our fitting and model selection procedures was characterised using synthetic data with realistic noise properties. This characterisation showed that the model selection procedure identifies order 0 profiles with good reliability – less than 0.1% of profiles being assigned higher order models. Anisotropic Gaussian profiles are also assigned order 2 models correctly most of the time. As expected, the rate of misclassification as order 0 increases with decreased anisotropy, but this rate becomes negligible when the difference between the largest and smallest eigenvalues exceeds about 40% of the mean diffusivity. The rate of misclassification of anisotropic Gaussian profiles at orders above 2 increases with increased anisotropy, and reaches 8% for the most anisotropic DTs that are typical in clinical human brain data. This is reasonable to expect, since highly anisotropic DTs contain very low ADC measurements in which the noise component is significant and causes greater deviation from the Gaussian profile.

Noisy bi-tensor profiles from mixed prolate DTs are detected as order 4 with reasonable reliability when the two principal directions are non-parallel. Misclassification at lower orders is around 3% when the two principal directions are orthogonal and the tissues are assumed to be

mixed in equal proportion. This rate increases significantly as the principal directions become more aligned and the proportions become less balanced. Misclassification at orders above 4 is around 0.5%.

Our procedure was applied to human brain data collected with parameters typical of those used in clinical scanners and appeared to classify isotropic and anisotropic regions correctly as order 0 and order 2, respectively. Several regions, in particular, the pons, the optic radiation and the corona radiation, were found consistently to contain a significant proportion of order 4 models. We confirmed that these findings are genuine, rather than an artefact of our procedure, by simulating the distribution of Gaussian models of the diffusion in these regions and showing that the corresponding ADC profiles are reliably classified as order 2.

It is clear from the results we have presented that higher order effects can be observed in clinical multi-direction DW-MR measurements taken from the brain, which render the standard DT model of the ADC profile inaccurate. In order to capture these effects, higher order models are required. When using higher order models, however, great care must be taken to choose the correct level of complexity so as not to discard genuine effects, while avoiding the random effects of noise, which can lead to false inferences from the data. Here we have provided the ground work for processing DW images in such a way as to detect and capture higher order effects reliably. However, there is considerable further work to be done in this area and we use the remainder of this discussion to focus on the strengths and weaknesses of the methods we have used and to highlight areas for future investigation.

In section 2, we described how the SH series can be fit efficiently and robustly via a linear transformation of the ADC samples. These functions thus provide a fast way to compute both the DT and higher order models for the ADC. The fitting and model selection procedures, as well as

the post-processing software used here are all written in Java and use only custom libraries. To fit the SH series up to order 2, takes approximately 12s per 128x128 slice of our images on a 700MHz PIII processor; up to order 8 requires approximately 90s per slice without model selection and 300s with. Further optimisation of both procedures is still possible, which should reduce these computation times. The SHs provide a convenient set of models for the ADC profile, but there may be alternative hierarchies that are more suitable for the kinds of profile we expect to observe given the underlying physical processes – for example, we might fit multi-tensor models directly. In addition to the computational simplicity of fitting SH models, however, an advantage of their use is the *avoidance* of adherence to a particular model of the underlying processes, which are not well understood. Where complex interactions of processes occur in the brain, for example, in the corona radiata, a multi-tensor model may be too restrictive to describe the ADC profile properly. Furthermore, the multi-tensor model assumes that the rate of exchange of particles between the tissue compartments is negligible, which may not be well justified.

We note here that the ADC is an inherently Gaussian quantity in as much as it only makes sense to take the logarithm of the ratio of the weighted and unweighted MR measurements, S/S_0 , when Eq. (5) holds. In general, a better base quantity, whose profile to model, might be the exponent of the ADC, i.e. S/S_0 itself, which relates more directly to the particle density function, [8]. Here, however, we have used the ADC throughout for consistency with familiar representations of diffusion data, as this does not affect our ability to detect Gaussian versus non-Gaussian behaviour.

In employing ANOVA and the F-test we implicitly assume that the errors in the ADC measurements are Gaussian. This is not true in general given that the measurements are obtained

from a log ratio of the original scanner measurements (see equation (1)), which follow Rice distributions, [22], although in practice the Gaussian provides a reasonable approximation. There are many other statistical tests that may be used in order to select the most appropriate from a set of candidate models and the use of alternative tests that explicitly take into account the expected noise may improve classification rates. Such a test should help overcome effects like the increased overfitting of order 2 profiles that occurs at high anisotropy, where higher levels of noise cause greater departure from the Gaussian profile, see section 3.1. Analysis of other techniques is the focus of current work, but the method used here has proved more effective than most of the common alternatives and has the advantage that, because it does not require an explicit noise model, it is simple and efficient to compute and generalises easily.

Almost all of the higher order models selected in our data are order 4 SH models, which suggests that the profiles arise from mixtures of two anisotropic tissues. When more than two such tissues intersect within a single voxel, we might expect that even higher order, e.g., 6 or 8, models to be required in order to model the profile successfully. The proportion of these models selected by the method described above is negligible in our human brain data, but we note that we have not explicitly tried to detect such models. Investigations on synthetic data suggest that in order to detect intersections of three or more WM tracts, represented by a tri-tensor model, with a reasonable true positive rate, the rate of overfitting of lower order models becomes unacceptable with our imaging parameter settings. For this reason, we set the order 4 versus higher order model decision threshold to the same value as that for order 2. In images acquired with a higher b-value, it is reasonable to expect that we may be able to detect these higher order profiles more reliably and a suitable decision threshold would need to be found in a similar way to that described in section 3.1.

There are many avenues for further work that lead on from the methods described in this paper. In our work so far, we have tested our methods on only data typical of clinical scanners ($b \approx 1000 \text{ s/mm}^2$) and indeed it is the purpose of this paper to show that non-Gaussian diffusion behaviour is significant in this data. However, our synthetic experiments show that with these imaging parameters, many regions in which higher order behaviour occurs can be missed. Alexander, et al, [12], showed that the kinds of higher order effects we observe are more pronounced with higher b , and one essential area for further work is to test our methods on data acquired with a range of b -values in order to determine the levels at which higher order effects become observable in different areas of the brain.

Analogues of some of the familiar scalar measures associated with the DT such as the mean diffusivity, anisotropy, etc, can be extracted from higher order models in a reasonably straightforward manner. These terms can be redefined slightly, as moments of the spherical functions described by our ADC profile models, as described in [25]. Frank, [16], noted that more stable measures of ADC profile shape are obtained by measuring the variation of the raw measurements rather than the DT model, however, indices computed on the raw measurements in this way are vulnerable to the effects of noise. The selection of a model of suitable complexity for each profile greatly reduces the overestimation of anisotropy that occurs as a result of computing the index on the raw measure and should improve the stability and accuracy of this kind of index. Full discussion is beyond the scope of this paper, however, and these indices will be described and investigated in detail elsewhere. Further work will be performed to investigate the stability of these measures with respect to varying imaging parameters and signal to noise.

We note that physical interpretation of higher order models of the ADC profile is more elusive than for simpler models, such as the DT, since they do not relate as directly to the particle

density function, which is indicative of the orientations of WM tracts – single or multiple. For example, the measurement acquired from the pons region, shown in Figure 9, is typical of the type of measurement that would arise from the mixture of vertical and horizontal fibres combined with the bi-tensor model, see Figure 1. The peaks in the profile, however, do not occur in these two directions, but rather in the intermediate directions. If we assume a model such as the bi-tensor model, the fibre directions can be obtained, but, as mentioned previously, such a model may not be appropriate in many cases and further work is necessary to investigate how to extract physiological information from the ADC profile in general.

Acknowledgements

Gareth Barker is funded by the Multiple Sclerosis Society of Great Britain and Northern Ireland. Thanks to Dr. Claudia Wheeler-Kingshott and Dr. Philip Boulby for pulse sequence development and to Dr. Olga Ciccarelli for providing data and aiding in the initial interpretation of the results.

References

1. P.J. Basser, J. Mattiello, and D. Le Bihan, *MR Diffusion Tensor Spectroscopy and Imaging*, Biophysical Journal, Vol. 66, 259-267, 1994.
2. P.J. Basser and C. Pierpaoli, *Microstructural and Physiological Features of Tissues Elucidated by Quantitative Diffusion Tensor MRI*, Journal of Magnetic Resonance, B. Vol. 111, pp. 209-219, 1996.
3. C. Pierpaoli, P. Jezzard, P.J. Basser, A. Barnett, G. Di Chiro, *Diffusion Tensor MR Imaging of the Human Brain*, Radiology, Vol. 201, No. 3, pp. 637-648, 1996.
4. J. Crank, *The Mathematics of Diffusion*, Oxford University Press, Oxford, UK, 1975.
5. E. O. Stejskal and J. E. Tanner, *Spin Diffusion Measurements: spin echoes in the presence of time-dependent field gradient*, J. Chem. Phys. Vol. 42, pp. 288-292, 1965.
6. D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis and M. Laval-Jeantet, *MR Imaging of Intravoxel Incoherent Motions: Application to Diffusion and Perfusion in Neurologic Disorders*, Radiology, Vol. 161, pp. 401-407, 1986.
7. D. K. Jones, M. A. Horsfield and A. Simmons, *Optimal Strategies for Measuring Diffusion in Anisotropic Systems by Magnetic Resonance Imaging*, Magnetic

- Resonance in Medicine, Vol. 42, pp. 515-525, 1999.
8. D. G. Cory and A. N. Garroway. *Measurement of translational displacement probabilities by NMR: an indicator of compartmentation*. Magnetic Resonance in Medicine, Vol. 14, pp. 435, 1990.
 9. D. Le Bihan, *Molecular diffusion, tissue microdynamics and microstructure*, NMR Biomed., Vol. 8, pp. 375-386, 1995.
 10. T. Niendorf, R. M. Dijkhuizen, D. G. Norris, C. M. van Lookeren, K. Nikolay. *Biexponential Diffusion Attenuation in Various States of Brain Tissue: Implications for Diffusion-Weighted Imaging*, Magnetic Resonance in Medicine, Vol. 36, pp. 847-857, 1996.
 11. J. Karger, H. Pfeifer and W. Heink. *Principles and applications of self diffusion measurements by nuclear magnetic resonance*, Advances in Magnetic Resonance, Vol. 12, pp. 1-89, 1988.
 12. A. L. Alexander, K. M. Hasan, M. Lazar, J. S. Tsuruda and D. L. Parker, *Analysis of Partial Volume Effects in Diffusion-Tensor MRI*, Magnetic Resonance in Medicine, Vol. 45, pp. 770-780, 2001.
 13. L. R. Frank. *Characterization of anisotropy in high angular resolution diffusion weighted MRI*. In: Proc 9th Annual Meeting ISMRM, Philadelphia, pp. 1531, 2001.
 14. D. S. Tuch, R. M. Weisskoff, J. W. Belliveau and V. J. Wedeen, *High Angular Resolution Diffusion Imaging of the Human Brain*. In: Proc 7th Annual Meeting ISMRM, Philadelphia, pp. 321, 1999.
 15. V. J. Wedeen, T. G. Reese, D. S. Tuch, M. R. Weigel, J-G. Dou, R. M. Weisskoff and D. Chesler. *Mapping fiber orientation spectra in cerebral white matter with Fourier transform diffusion MRI*. In: Proc 8th Annual Meeting ISMRM, Denver, pp. 82, 2000.
 16. L. R. Frank, *Anisotropy in High Angular Resolution Diffusion-Weighted MRI*, Magnetic Resonance in Medicine, Vol. 45, pp. 935-939, 2001.
 17. B. A. Inglis, E. L. Bossart, D. L. Buckley, E. D. Wirth III and T. H. Mareci, *Visualization of Neural Tissue Water Compartments Using Biexponential Diffusion Tensor MRI*, Magnetic Resonance in Medicine, Vol. 45, pp. 580-587, 2001.
 18. A. F. Nikiforov and V. B. Uvarov, *Special Functions of Mathematical Physics*, Birkhauser Verlag Basel, 1988.
 19. W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical Recipes in C*, Cambridge University Press, 1988.
 20. C. H. Brechbuhler, G. Gerig and O. Kubler, *Parametrization of Closed Surfaces for 3D Shape Description*, Computer Vision and Image Understanding, Vol. 61, pp. 154-170, 1995.
 21. P. Armitage and G. Berry, *Statistical Methods in Medical Research*, Blackwell Scientific Publications, 1971.
 22. A. J. Miller and P. M. Joseph, *The use of power images to perform quantitative*

- analysis on low SNR MR images*. Magnetic Resonance Imaging, Vol. 11: 1051-1056, 1993.
23. S. Pajevic and C. Pierpaoli, *Color Schemes to Represent the Orientation of Anisotropic Tissues from Diffusion Tensor Data: Application to White Matter Fiber Tract Mapping in the Human Brain*. Magnetic Resonance Medicine, Vol. 42, pp. 526-540, 1999.
 24. P. A. Armitage and M. E. Bastin, *Utilizing the Diffusion-to-noise ratio to Optimize Magnetic Resonance Diffusion Tensor Acquisition Strategies for Improving Measurements of Diffusion Anisotropy*. Magnetic Resonance in Medicine, Vol. 45, pp. 1056-1065, 2001.
 25. D. C. Alexander, *Note on Indices of Shape and Similarity for Diffusion Tensors*, UCL, Department of Computer Science Research Note, RN/01/14, 2001.

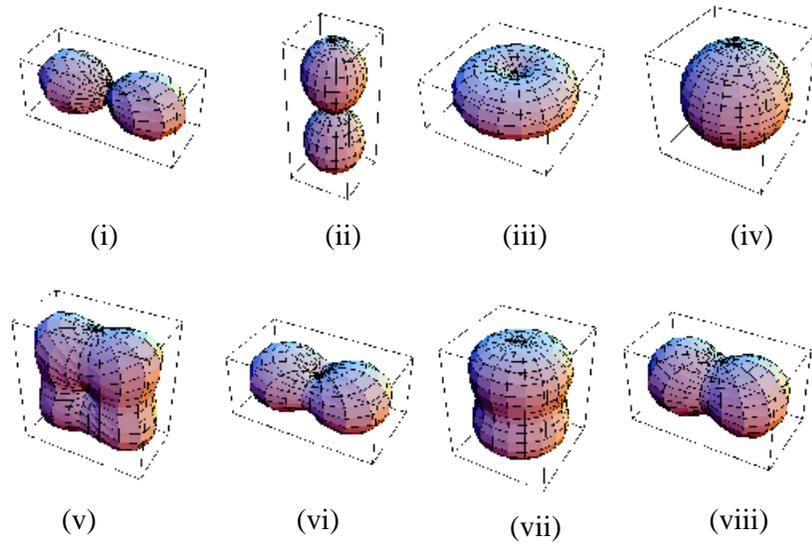


Figure 1 Simulated examples of ADC profiles. Top row: profiles corresponding to Gaussian diffusion processes; (i) prolate DT oriented along x-axis (eigenvalues $[1700, 200, 200] \times 10^{-6} \text{ mm}^2/\text{s}$), (ii) prolate DT oriented along z-axis (eigenvalues $[200, 200, 1700] \times 10^{-6} \text{ mm}^2/\text{s}$), (iii) oblate DT (eigenvalues $[950, 950, 200] \times 10^{-6} \text{ mm}^2/\text{s}$), (iv) isotropic DT (eigenvalues $[700, 700, 700] \times 10^{-6} \text{ mm}^2/\text{s}$). Bottom row: ADC profiles corresponding to the bi-tensor model combining pairs of DTs from the top row in equal proportion with b set to $1000 \text{ s}/\text{mm}^2$; (v) combines (i) and (ii), (vi) combines (i) and (iii), (vii) combines (ii) and (iii), and (viii) combines (i) and (iv).

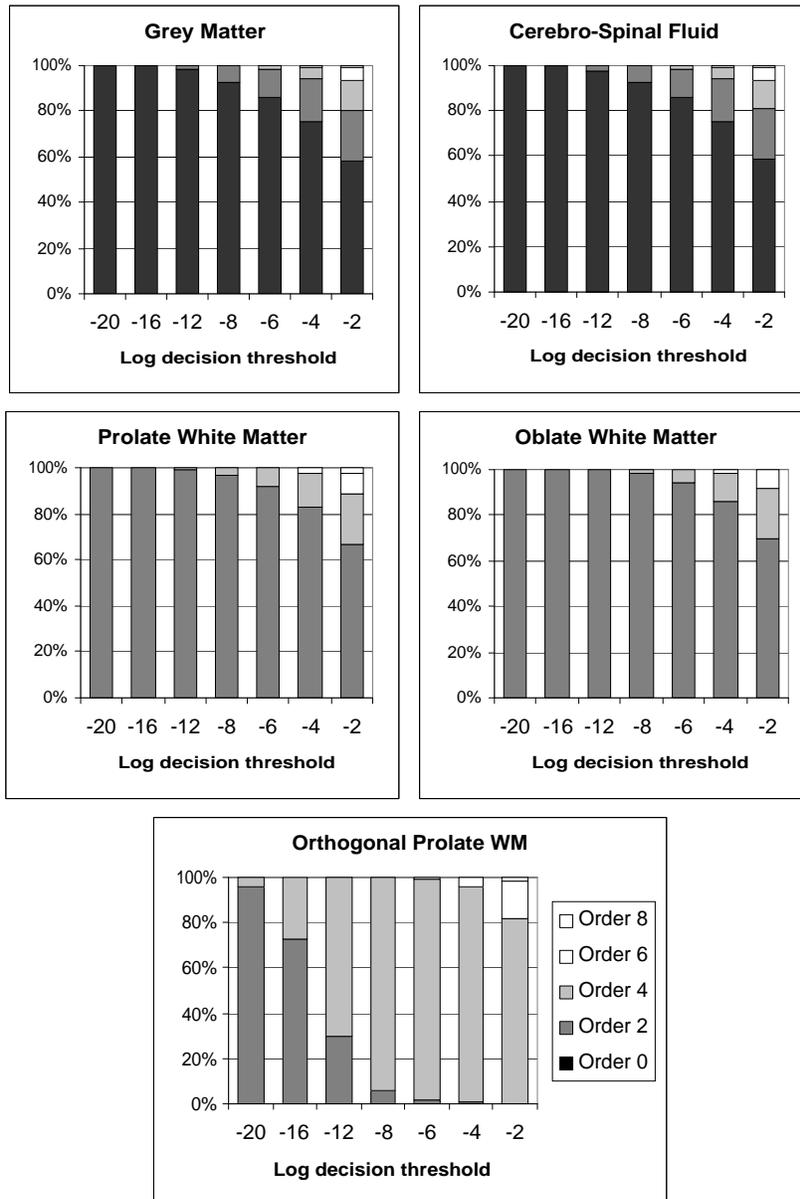


Figure 2 Proportions of SH series model orders selected for a variety (see text) of ADC profiles from simulated measurements with varying F-test decision threshold.

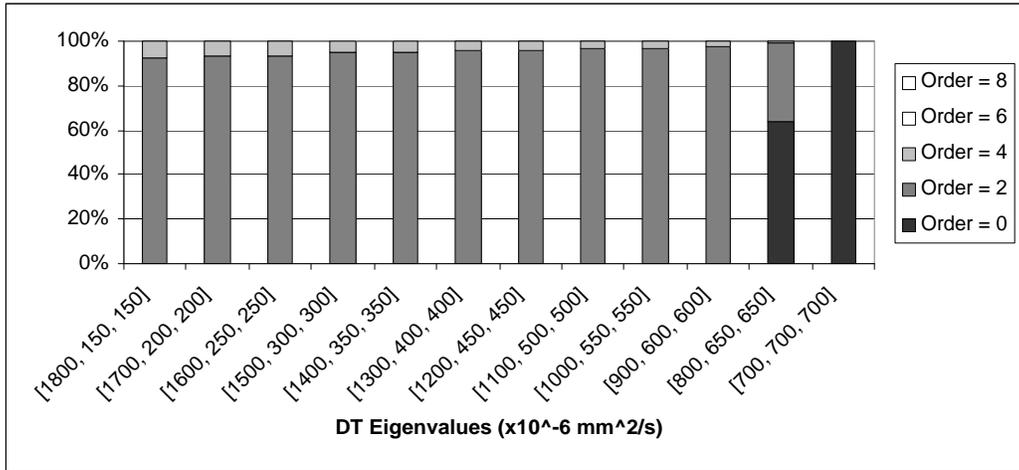


Figure 3 Proportions of SH series model orders selected for profiles of simulated measurements from prolate DTs with varying levels of anisotropy (increasing right to left).

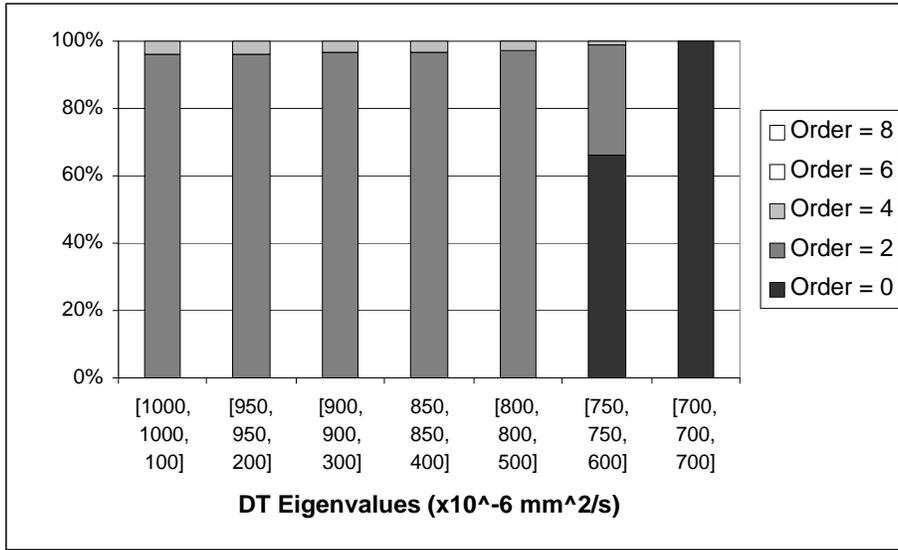


Figure 4 Proportions of SH series model orders selected for profiles of simulated measurements from oblate DTs with varying levels of anisotropy (increasing right to left).

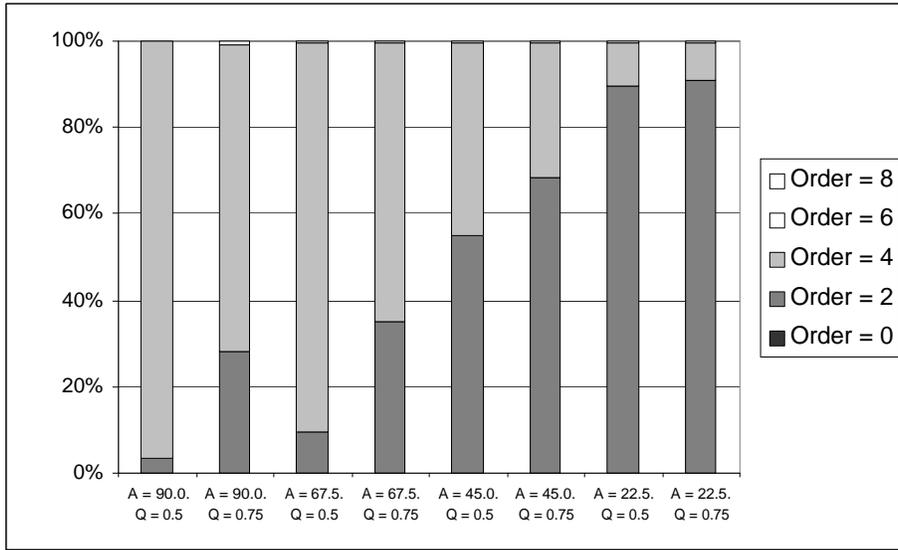


Figure 5 Proportions of SH series model orders selected for profiles of simulated measurements of a bi-tensor ADC profile from a combination of two prolate DTs intersecting at various angles, A. Results are included for cases where there are equal proportions ($Q = 0.5$) of the two tissue types and where there is 3 times more of one type ($Q = 0.75$).

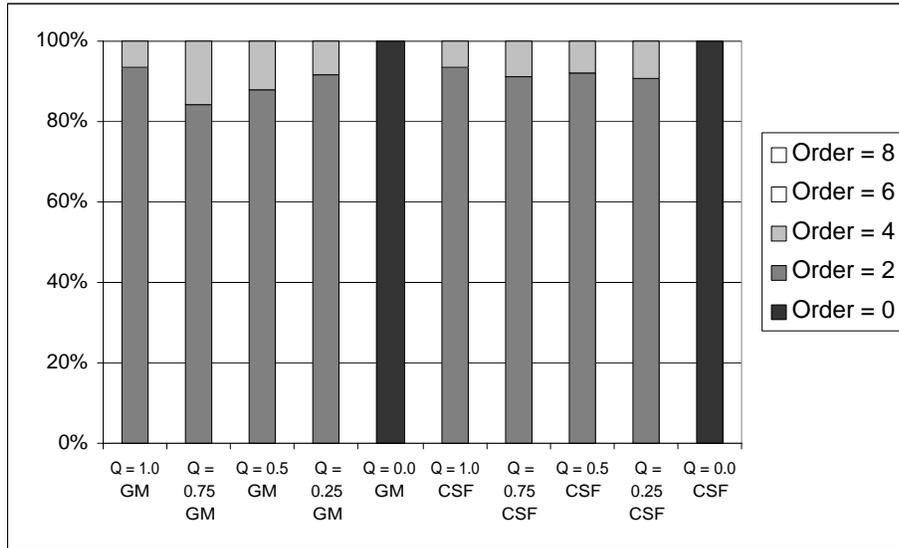


Figure 6 Proportions of SH series model orders selected for profiles of simulated measurements of bi-tensor ADC profiles from a combination of a prolate DT and isotropic DTs emulating those found in grey matter and CSF.

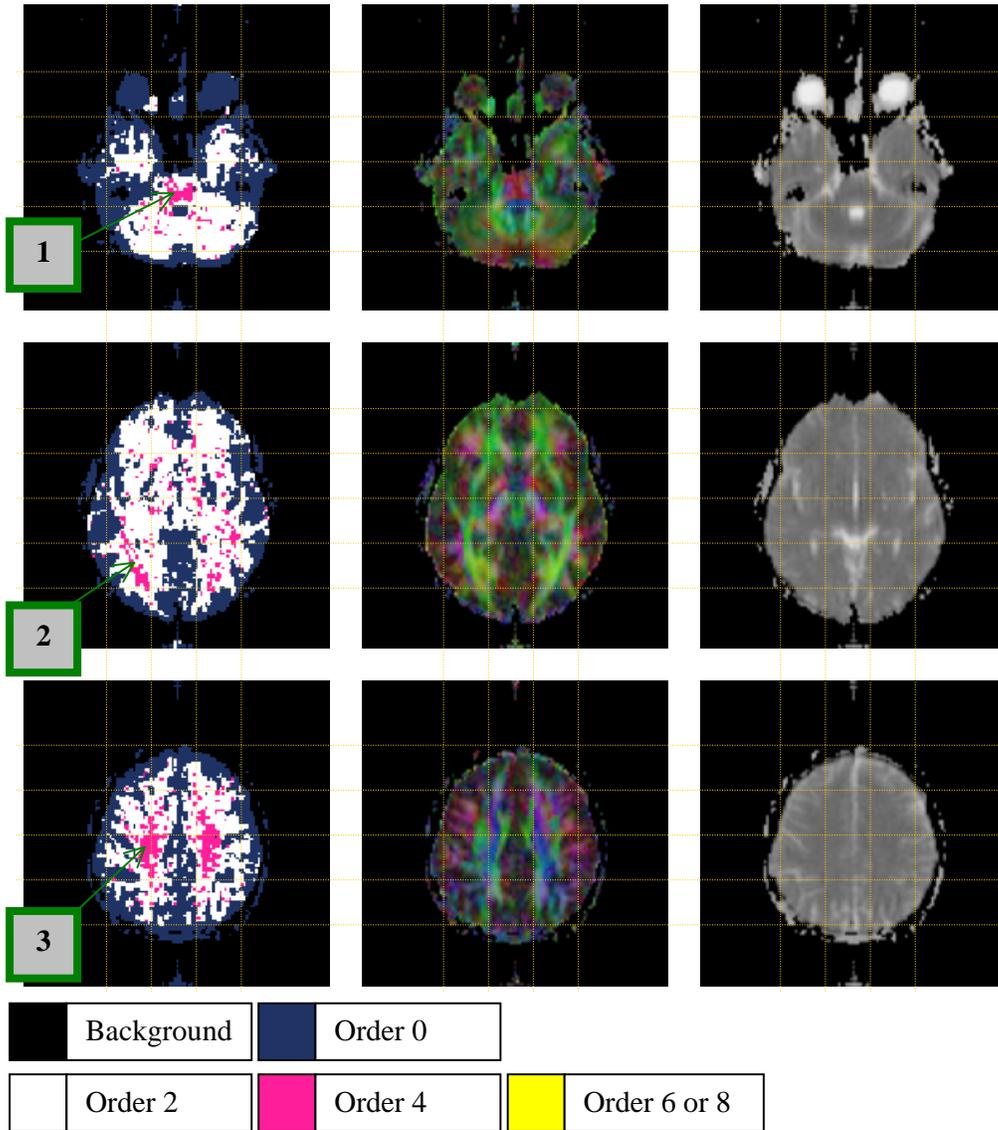


Figure 7 Order maps from human brain data (left) together with colour coded principal DT direction maps (middle) and mean diffusivity maps (right). The key below the images refers to the order maps in the left hand column.

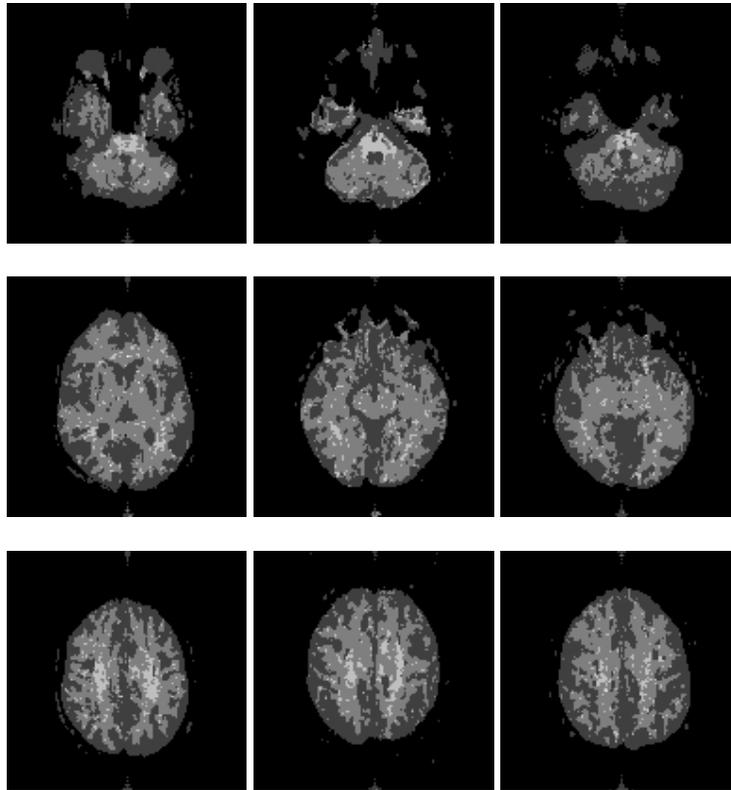


Figure 8 Slices of the orders maps computed for the three data sets not shown in Figure 7 through the pons, the optic radiation and the corona radiata at similar levels to those shown in Figure 7. The data sets shown in the left and centre columns exhibit strong clusters in the same anatomic regions as those highlighted on Figure 7. The data set on the right is of somewhat poorer quality and, although these clusters still exist, they are less pronounced.

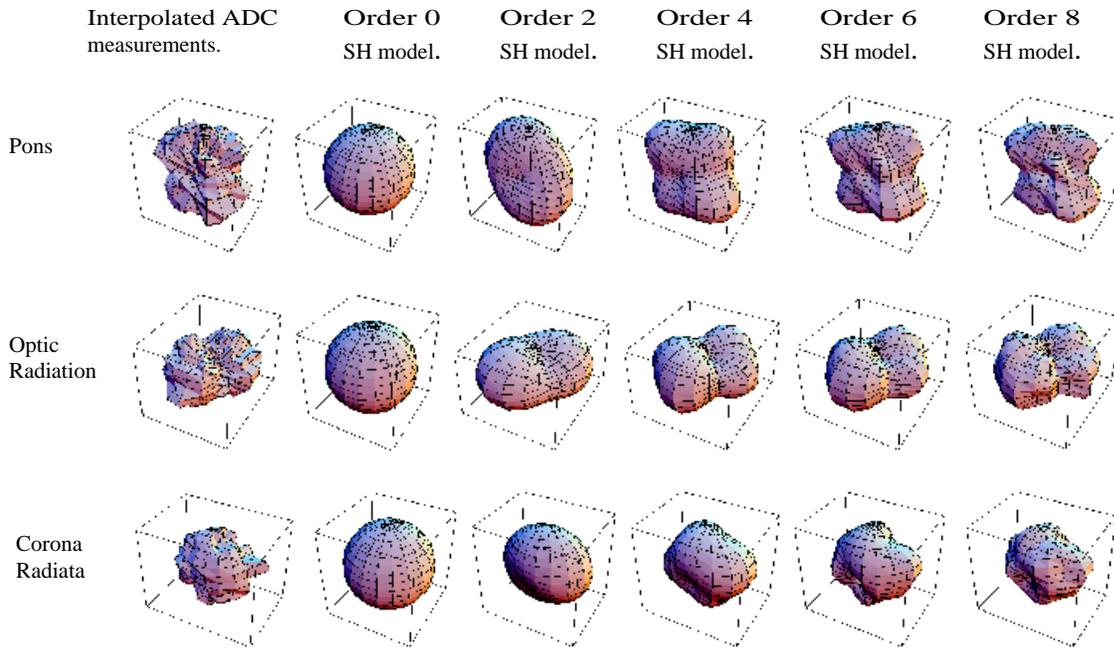


Figure 9 Typical measurements (left) together with SH models of orders 0, 2, 4, 6, 8 (second from left to right) from each of the three higher order clusters labelled in Figure 7. Top: measurement from pons (order 4 model selected). Middle: measurement from optic radiation (order 4 model selected). Bottom: measurement from corona radiata (order 4 model selected).

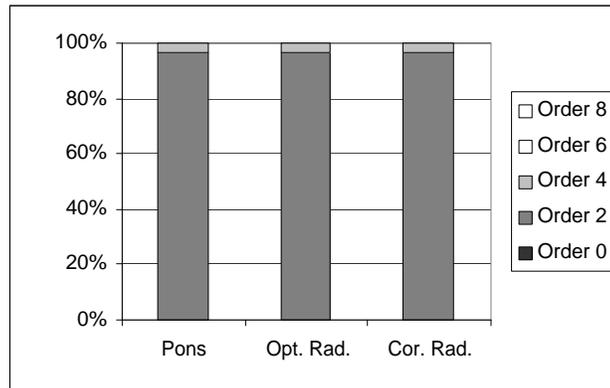


Figure 10 Proportions of SH series model orders selected for simulated ADC profiles drawn from the distributions of DT parameters from three higher order regions of the human brain.

This document was created with Win2PDF available at <http://www.daneprairie.com>.
The unregistered version of Win2PDF is for evaluation or non-commercial use only.